

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) EP 1 096 751 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
02.05.2001 Bulletin 2001/18

(51) Int. Cl.⁷: H04L 29/06

(21) Application number: 00203642.4

(22) Date of filing: 18.10.2000

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(72) Inventors:
• Iterum, Skef F.
San Francisco, CA 94107 (US)
• Murphy, Declan J.
San Francisco, CA 94131 (US)

(30) Priority: 21.10.1999 US 160992 P
15.09.2000 US 662553

(74) Representative:
Hanna, Peter William Derek
Peter Hanna Associates
11 Mespil Road
Dublin 4 (IE)

(71) Applicant:
SUN MICROSYSTEMS, INC.
Palo Alto, California 94303 (US)

(54) Method and apparatus for reaching agreement between nodes in a distributed system

(57) One embodiment of the present invention provides a system for selecting a node to host a primary server for a service from a plurality of nodes in a distributed computing system. The system operates by receiving an indication that a state of the distributed computing system has changed. In response to this indication, the system determines if there is already a node hosting the primary server for the service. If not, the system selects a node to host the primary server using the assumption that a given node from the plurality of nodes in the distributed computing system hosts the primary server. The system then communicates rank information between the given node and other nodes in the distributed computing system, wherein each node in the distributed computing system has a unique rank with respect to the other nodes in the distributed computing system. The system next compares the rank of the given node with the rank of the other nodes in the distributed computing system. If one of the other nodes has a higher rank than the given node, the system disqualifies the given node from hosting the primary server.

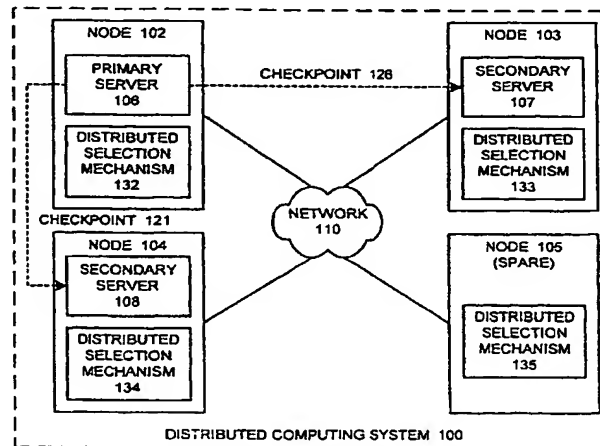


FIG. 1

EP 1 096 751 A2

Description

Related Application

[0001] This application hereby claims priority under 5
35 U.S.C. § 119 to U. S. Provisional Patent Application
No. 60/160,992 filed on October 21, 1999, entitled "Dis-
tributed Multi-Tier Mechanism for Agreement."

BACKGROUND

Field of the Invention

[0002] The present invention relates to coordinating
activities between nodes in a distributed computing sys-
tem. More specifically, the present invention relates to a
method and an apparatus for reaching agreement
between nodes in the distributed computing system
regarding a node to function as a primary provider for a
service.

Related Art

[0003] As computer networks are increasingly used
to link computer systems together, distributed comput-
ing systems have been developed to control interactions
between computer systems. Some distributed comput-
ing systems allow client computer systems to access
resources on server computer systems. For example, a
client computer system may be able to access informa-
tion contained in a database on a server computer sys-
tem.

[0004] When a server computer system fails, it is
desirable for the distributed computing system to auto-
matically recover from this failure. Distributed computer
systems possessing an ability to recover from such
server failures are referred to as "highly available sys-
tems."

[0005] For a highly available system to function
properly, the highly available system must be able to
detect a server failure and reconfigure itself so that
accesses to a failed server are redirected to a backup
secondary server.

[0006] One problem in designing such a highly
available system is that some distributed computing
system functions must be centralized in order to operate
efficiently. For example, it is desirable to centralize an
arbiter that keeps track of where primary and secondary
copies of a server are located in a distributed computing
system. However, a node that hosts such a centralized
arbiter may itself fail. Hence, it is necessary to provide a
mechanism to select a new node to host the centralized
arbiter.

[0007] Moreover, this selection mechanism must
operate in a distributed fashion because, for the rea-
sons stated above, no centralized mechanism is certain
to continue functioning. Furthermore, it is necessary for
the node selection process to operate so that the nodes

that remain functioning in the distributed computing sys-
tem agree on the same node to host the centralized
arbiter. For efficiency reasons, it is also desirable for the
node selection mechanism not to move the centralized
arbiter unless it is necessary to do so.

[0008] Hence, what is needed is a method and an
apparatus that operates in a distributed manner to
select a node to host a primary server for a service.

10 SUMMARY

[0009] One embodiment of the present invention
provides a system for selecting a node to host a primary
server for a service from a plurality of nodes in a distrib-
uted computing system. The system operates by receiv-
ing an indication that a state of the distributed
computing system has changed. In response to this
indication, the system determines if there is already a
node hosting the primary server for the service. If not,
the system selects a node to host the primary server
using the assumption that a given node from the plural-
ity of nodes in the distributed computing system hosts
the primary server. The system then communicates
rank information between the given node and other
nodes in the distributed computing system, wherein
each node in the distributed computing system has a
unique rank with respect to the other nodes in the dis-
tributed computing system. The system next compares
the rank of the given node with the rank of the other
nodes in the distributed computing system. If one of the
other nodes has a higher rank than the given node, the
system disqualifies the given node from hosting the pri-
mary server.

[0010] In one embodiment of the present invention,
if there exists a node to host the primary server, the sys-
tem allows the node that hosts the primary server to
communicate with other nodes in the distributed com-
puting system in order to disqualify the other nodes from
hosting the primary server.

[0011] In one embodiment of the present invention,
the system maintains a candidate variable in the given
node identifying a candidate node to host the primary
server. In a variation on this embodiment, the system
initially sets the candidate variable to identify the given
node.

[0012] In one embodiment of the present invention,
after a new node has been selected to host the primary
server, if the new node is different from a previous node
that hosted the primary server, the system maps con-
nections for the service to the new node. In a variation
on this embodiment, the system also configures the
new node to host the primary server for the service.

[0013] In one embodiment of the present invention,
the system restarts the service if the service was inter-
rupted as a result of the change in state of the distrib-
uted computing system.

[0014] In one embodiment of the present invention,
the given node in the distributed computing system can

act as one of: a host for the primary server for the service; a host for a secondary server for the service, wherein the secondary server periodically receives checkpointing information from the primary server; or a spare for the primary server, wherein the spare does not receive checkpointing information from the primary server.

[0015] In one embodiment of the present invention, upon initial startup of the service, the system selects a highest ranking spare to host the primary server for the service.

[0016] In one embodiment of the present invention, the system allows the primary server to configure spares in the distributed computing system to host secondary servers for the service.

[0017] In one embodiment of the present invention, comparing the rank of the given node with the rank of the other nodes in the distributed computing system involves considering a host for a secondary server to have a higher rank than a spare.

[0018] In one embodiment of the present invention, after disqualifying the given node from hosting the primary server, the system ceases to communicate rank information between the given node and the other nodes in the distributed computing system.

BRIEF DESCRIPTION OF THE FIGURES

[0019]

FIG. 1 illustrates a distributed computing system in accordance with an embodiment of the present invention.

FIG. 2 illustrates how highly available services are controlled within a distributed computing system in accordance with an embodiment of the present invention.

FIG. 3 illustrates how a replica managers controls highly available services in accordance with an embodiment of the present invention.

FIG. 4 is a flow chart illustrating the process of selecting and configuring a new primary server in accordance with an embodiment of the present invention.

FIG. 5 is a flow chart illustrating some of the operations performed by a primary server in accordance with an embodiment of the present invention.

FIG. 6 illustrates how a node is selected to host a primary server through a disqualification process in accordance with an embodiment of the present invention.

FIG. 7 illustrates how nodes a disqualified in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

[0020] The following description is presented to

enable any person skilled in the art to make and use the invention, and is provided in the context of a particular application and its requirements. Various modifications to the disclosed embodiments will be readily apparent to those skilled in the art, and the general principles defined herein may be applied to other embodiments and applications without departing from the spirit and scope of the present invention. Thus, the present invention is not intended to be limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features disclosed herein.

[0021] The data structures and code described in this detailed description are typically stored on a computer readable storage medium, which may be any device or medium that can store code and/or data for use by a computer system. This includes, but is not limited to, magnetic and optical storage devices such as disk drives, magnetic tape, CDs (compact discs) and DVDs (digital versatile discs or digital video discs), and computer instruction signals embodied in a transmission medium (with or without a carrier wave upon which the signals are modulated). For example, the transmission medium may include a communications network, such as the Internet.

Distributed Computing System

[0022] FIG. 1 illustrates a distributed computing system 100 in accordance with an embodiment of the present invention. Distributed computing system 100 includes a number of computing nodes 102-105, which are coupled together through a network 110.

[0023] Network 110 can include any type of wire or wireless communication channel capable of coupling together computing nodes. This includes, but is not limited to, a local area network, a wide area network, or a combination of networks. In one embodiment of the present invention, network 110 includes the Internet. In another embodiment of the present invention, network 110 is a local high speed network that enables distributed computing system 100 to function as a clustered computing system (hereinafter referred to as a "cluster").

[0024] Nodes 102-105 can generally include any type of computer system, including, but not limited to, a computer system based on a microprocessor, a main-frame computer, a digital signal processor, a personal organizer, a device controller, and a computational engine within an appliance.

[0025] Nodes 102-105 also host servers, which include a mechanism for servicing requests from a client for computational and/or data storage resources. More specifically, node 102 hosts primary server 106, which services requests from clients (not shown) for a service involving computational and/or data storage resources.

[0026] Nodes 103-104 host secondary servers 107-108, respectively, for the same service. These second-

ary servers act as backup servers for primary server 106. To this end, secondary servers 107-108 receive periodic checkpoints 120-121 from primary server 106. These periodic checkpoints enable secondary servers 107-108 to maintain consistent state with primary server 106. This makes it possible for one of secondary servers 107-108 to take over for primary server 106 if primary server 106 fails.

[0027] Node 105 can serve as a spare node to host the service provided by primary server 106. Hence, node 105 can be configured to host a secondary server with respect to a service provided by primary server 106. Alternatively, if all primary servers and secondary servers for the service fail, node 105 can be configured to host a new primary server for the service.

[0028] Also note that nodes 102-105 contain distributed selection mechanisms 132-135, respectively. Distributed selection mechanisms 132-135 communicate with each other to select a new node to host primary server 106 when node 102 fails or otherwise becomes unavailable. This process is described in more detail below with reference to FIGs. 2-6.

Controlling Highly Available Services

[0029] FIG. 2 illustrates how highly available services 202-205 are controlled within distributed computing system 100 in accordance with an embodiment of the present invention. Note that highly available services 202-205 continue to operate even if individual nodes of distributed computing system 100 fail.

[0030] Highly available services 202-205 operate under control of replica manager 206. Referring to FIG. 3, for each service, replica manager 206 keeps a record of which nodes in distributed computing system 100 function as primary servers, and which nodes function as secondary servers. For example, in FIG. 3 replica manager 206 keeps track of highly available services 202-205. The primary server for service 202 is node 103, and the secondary servers are nodes 104 and 105. The primary server for service 203 is node 104, and the secondary servers are nodes 103 and 105. The primary server for service 204 is node 102, and the secondary servers are nodes 104-105. The primary server for service 205 is node 103, and the secondary servers are nodes 102, 104 and 105.

[0031] Replica manager 206 additionally performs a number of related functions, such as configuring a node to host a primary (which may involve demoting a current host for the primary to host a secondary). Replica manager 206 may additionally perform other functions, such as: adding a service; removing providers for a service; registering providers for a service; removing a service; handling provider failures; bringing up new providers for a service; and handling dependencies between services (which may involve ensuring that primaries for dependent services are co-located on the same node).

[0032] Referring back to FIG. 2, replica manager 206 is itself a highly available service that operates under control of replica manager manager (RMM) 208. Note that RMM 208 is not managed by a higher level service.

[0033] As illustrated in FIG. 2, RMM 208 communicates with cluster membership monitor (CMM) 210. CMM 210 monitors cluster membership and alerts RMM 208 if any changes in the cluster membership occur.

[0034] CMM 210 uses transport layer 212 to exchange messages between nodes 102-105.

Process of Selecting a New Primary

[0035] FIG. 4 is a flow chart illustrating the process of selecting and configuring a primary server in accordance with an embodiment of the present invention. Note that this process is run concurrently by each active node in distributed computing system 100. The system begins by receiving an indication from CMM 210 that the membership in the cluster has changed (step 401).

[0036] In response to this indication, the system obtains a lock on a local candidate variable which contains an identifier for a candidate node to host primary server 106 (step 402). The system also obtains an additional lock to hold off requesters for the service (step 404).

[0037] Next, the system executes a disqualification process by communicating with other nodes in distributed computing system 100 in order to disqualify the other nodes from acting as the primary server 106 (step 406). This process is described in more detail with reference to FIG. 6 below.

[0038] After the disqualification process, the remaining node, which is not disqualified, becomes the primary node. If the node hosting primary server 106 has changed, this may involve re-mapping connections for the service to point to the new node (step 408). It may also involve initializing the new node to act as the host for the primary (step 410).

[0039] Finally, the service is started (step 412). This may involve unfreezing the service if it was previously frozen, as well as releasing the previously obtained lock that holds off requesters for the service.

[0040] FIG. 5 is a flow chart illustrating some of the operations performed by a primary server 106 in accordance with an embodiment of the present invention. During operation, primary server 106 performs periodic checkpointing operations 120-121 with secondary servers 107-108, respectively (step 502). These checkpointing operations allow secondary servers 107-108 to take over from primary server 106 if primary server 106 fails. Primary server 106 also periodically attempts to promote spare nodes (such as node 105 in FIG. 1) to host secondaries (step 504). This promotion process involves transferring state information to a spare node in order to bring the spare node up to date

with respect to the existing secondaries.

[0041] FIG. 6 illustrates how a node is selected to host primary server 106 through a disqualification process in accordance with an embodiment of the present invention. Note that FIG. 6 describes in more detail the process described above with reference to step 406 in FIG. 4.

[0042] The system starts by determining if a node that was previously hosting primary server 106 continues to exist (step 602).

[0043] If not, the system retrieves the state of a local provider for the service (step 604). The system then sets the candidate variable to identify the local provider (step 606), and subsequently unlocks the candidate lock that was set previously in step 402 of FIG. 4 (step 608). Next, if the candidate is not the local provider, the system ends the process (step 610).

[0044] Next, for all other nodes *l* in the cluster, the system attempts to disqualify node *l* by writing a new identifier into the candidate variable for node *l* if the rank of node *l* is less than the rank of the present node (step 612). This process is described in more detail with reference to FIG. 7 below. Finally, if node *l*'s local provider has a higher rank than the present node, the process terminates because the present node is disqualified (step 614).

[0045] Note that a rank of a node can be obtained by comparing a unique identifier for the node with unique identifiers for other nodes. Also note that the rank of a primary server is greater than the rank of a secondary server, and that the rank of a secondary server is greater than the rank of a spare. The above-listed restrictions on rank ensure that an existing primary that has not failed continues to function as the primary, and that an existing secondary will be chosen ahead of a spare. Of course, when the system is initialized, no primaries or secondaries exist, so a spare is selected to be the primary.

[0046] On the other hand, if the node that was hosting primary server 106 continues to function, the system sets the candidate to be this node (step 616), and unlocks the candidate node (step 618).

[0047] If the present node does not host primary server 106, the process is finished. Otherwise, if the present node is hosting primary server 106, the system considers each other node *l* in the cluster. If the present node has already communicated with *l*, the system skips node *l* (step 622). Otherwise, the system communicates with node *l* in order to disqualify node *l* from acting as the host for primary server 106 (step 624). This may involve causing an identifier for the present node to be written into the candidate variable for node *l*.

[0048] FIG. 7 illustrates how nodes are disqualified in accordance with an embodiment of the present invention. Note that FIG. 7 describes in more detail the process described above with reference to step 612 in FIG. 6. The caller first locks the candidate variable for node *l* (step 702). If the caller determines that the caller's pro-

vider has a higher rank than is specified in the candidate variable for *l*, the caller overwrites the candidate variable for *l* with the caller's provider (step 704). Next, the caller unlocks the candidate variable for *l* (step 706).

[0049] The foregoing descriptions of embodiments of the invention have been presented for purposes of illustration and description only. They are not intended to be exhaustive or to limit the present invention to the forms disclosed. Accordingly, many modifications and variations will be apparent to practitioners skilled in the art. Additionally, the above disclosure is not intended to limit the present invention. The scope of the present invention is defined by the appended claims.

15 Claims

1. A method for selecting a node 102,103,104,105 to host a primary server 106 for a service from a plurality of nodes in a distributed computing system 100, the method comprising:

- a) receiving 401 an indication that a state of the distributed computing system has changed;
- b) in response to the indication, determining 602 if there is already a node hosting the primary server for the service; and
- c) if there is not already a node hosting the primary server, selecting 604 a node to host the primary server based upon rank information for the nodes.

2. The method of claim 1, wherein selecting the node to host the primary server involves:

- a) assuming that a given node from the plurality of nodes in the distributed computing system hosts the primary server,
- b) communicating 612 rank information between the given node and other nodes in the distributed computing system, wherein each node in the distributed computing system has a unique rank with respect to the other nodes in the distributed computing system,
- c) comparing a rank of the given node with a rank of the other nodes in the distributed computing system, and
- d) if one of the other nodes in the distributed computing system has a higher rank than the given node, disqualifying 614 the given node from hosting the primary server.

3. The method of claim 2, further comprising, if there exists a node that is configured to host the primary server, allowing 622 the node that is configured to host the primary server to communicate with other nodes in the distributed computing system in order to disqualify 624 the other nodes from hosting the primary server.

4. The method of claim 2, wherein assuming that the given node hosts the primary server involves:
 - a) maintaining a candidate variable in the given node identifying a candidate node to host the primary server; and
 - b) initially setting the candidate variable to identify the given node.
5. The method of claim 1, further comprising, after a new node has been selected to host the primary server, if the new node is different from a previous node that hosted the primary server, establishing 408 connections for the service to the new node.
6. The method of claim 1, further comprising, after a new node has been selected to host the primary server, if the new node is different from a previous node that hosted the primary server, configuring the new node to host the primary server for the service.
7. The method of claim 1, further comprising restarting the service if the service was interrupted as a result of the change in state of the distributed computing system.
8. The method of claim 2, wherein the given node in the distributed computing system acts a one of:
 - a) a host for the primary server 106 for the service;
 - b) a host for a secondary server 107, 108 for the service, wherein the secondary server periodically receives checkpointing information from the primary server; and
 - c) a spare for the primary server, wherein the spare does not receive checkpointing information from the primary server.
9. The method of claim 8, further comprising, upon initial startup of the service, selecting a highest ranking spare to host the primary server for the service.
10. The method of claim 8, further comprising allowing the primary server to configure 504 spares in the distributed computing system to host secondary servers for the service.
11. The method of claim 8, wherein comparing the rank of the given node with the rank of the other nodes in the distributed computing system involves considering a host for the primary server to have a higher rank than a host for a space, and considering a host for a secondary server to have a higher rank than a spare.
12. The method of claim 2, wherein disqualifying the given node from hosting the primary server involves ceasing to communicate rank information between the given node and the other nodes in the distributed computing system.
13. A computer-readable storage medium storing instructions that when executed by a computer cause the computer to perform the method steps of any one of claims 1 to 12.
14. A computer program, which when run on a computer, is adapted to perform the method steps of any one of claims 1 to 12.
15. An apparatus that selects a node to host a primary server for a service from a plurality of nodes in a distributed computing system, the apparatus comprising:
 - a) a receiving mechanism 401 that is configured to receive an indication that a state of the distributed computing system has changed;
 - b) a determination mechanism 602 that is configured to determine if there is already a node hosting the primary server for the service in response to the indication;
 - c) a selecting mechanism 604, wherein if there is not already a node hosting the primary server, the selecting mechanism is configured to select a node to host the primary server based upon rank information for the nodes.
16. The apparatus of claim 15, wherein, in selecting a node to host the primary server based upon rank information, the selecting mechanism is configured to:
 - a) communicate rank information between the given node and other nodes in the distributed computing system, wherein each node in the distributed computing system has a unique rank with respect to the other nodes in the distributed computing system, and to
 - b) compare a rank of the given node with a rank of the other nodes in the distributed computing system.
17. The apparatus of claim 16, further comprising a disqualification mechanism that is configured to disqualify the given node from hosting the primary server if one of the other nodes in the distributed computing system has a higher rank than the given node.
18. The apparatus of claim 16, further comprising a mechanism on the primary server that is configured to communicate with other nodes in the distributed computing system in order to disqualify the other

nodes from hosting the primary server.

19. The apparatus of claim 16, wherein the selecting mechanism is configured to:

- a) maintain a candidate variable in the given node identifying a candidate node to host the primary server; and to
- b) initially set the candidate variable to identify the given node.

20. The apparatus of claim 15, further comprising a connection mechanism that is configured to establish connections for the service to a new node after the new node has been selected to host the primary server, and if the new node is different from a previous node that hosted the primary server.

21. The apparatus of claim 15, further comprising a mechanism that configures a new node to host the primary server for the service, after the new node has been selected to host the primary server, and if the new node is different from a previous node that hosted the primary server.

22. The apparatus of claim 15, further comprising a restarting mechanism that is configured to restart the service if the service was interrupted as a result of the change in state of the distributed computing system.

23. The apparatus of claim 16, wherein the given node in the distributed computing system acts a one of:

- a) a host for the primary server for the service;
- b) a host for a secondary server for the service, wherein the secondary server periodically receives checkpointing information from the primary server; and
- c) a spare for the primary server, wherein the spare does not receive checkpointing information from the primary server.

24. The apparatus of claim 23, further comprising an initialization mechanism wherein during initialization of the service, the initialization mechanism is configured to select a highest ranking spare to host the primary server for the service.

25. The apparatus of claim 23, further comprising a promotion mechanism on the primary server that is configured to promote spares in the distributed computing system to host secondary servers for the service.

26. The apparatus of claim 23, wherein while comparing the rank of the given node with the rank of the other nodes in the distributed computing system,

the selecting mechanism is configured to consider a host for the primary server to have a higher rank than a host for a secondary server, and to consider a host for a secondary server to have a higher rank than a spare.

27. The apparatus of claim 16, wherein the selecting mechanism is configured to cease to communicate rank information between the given node and the other nodes in the distributed computing system after the given node is disqualified by the disqualification mechanism.

28. A method for selecting a node to host a primary server for a service from a plurality of nodes in a distributed computer system, comprising:

- a) communicating disqualification information between the node and remaining nodes in the plurality of nodes;
- b) disqualifying the node from hosting the primary server based upon the disqualification information received from the remaining nodes.

29. The method of claim 28, wherein the disqualification information comprises a node rank information.

30. The method of claim 29, wherein the node rank for a given node is calculated using an assumption that the given node hosts the primary server.

31. The method of claim 30, wherein the calculated node rank is unique with respect to the ranks of other nodes in the distributed computer system.

32. The method of claim 29, wherein the disqualifying of the node comprises:

- a) comparing a rank of the node to a set of ranks of the remaining nodes in the distributed computer system; and
- b) disqualifying the node from hosting the primary server if one of the set of ranks of the remaining nodes is higher than the rank of the node.

33. The method of claim 28, further comprising repeating the acts of communicating disqualification information and disqualifying the node for at least one more node in the plurality of nodes.

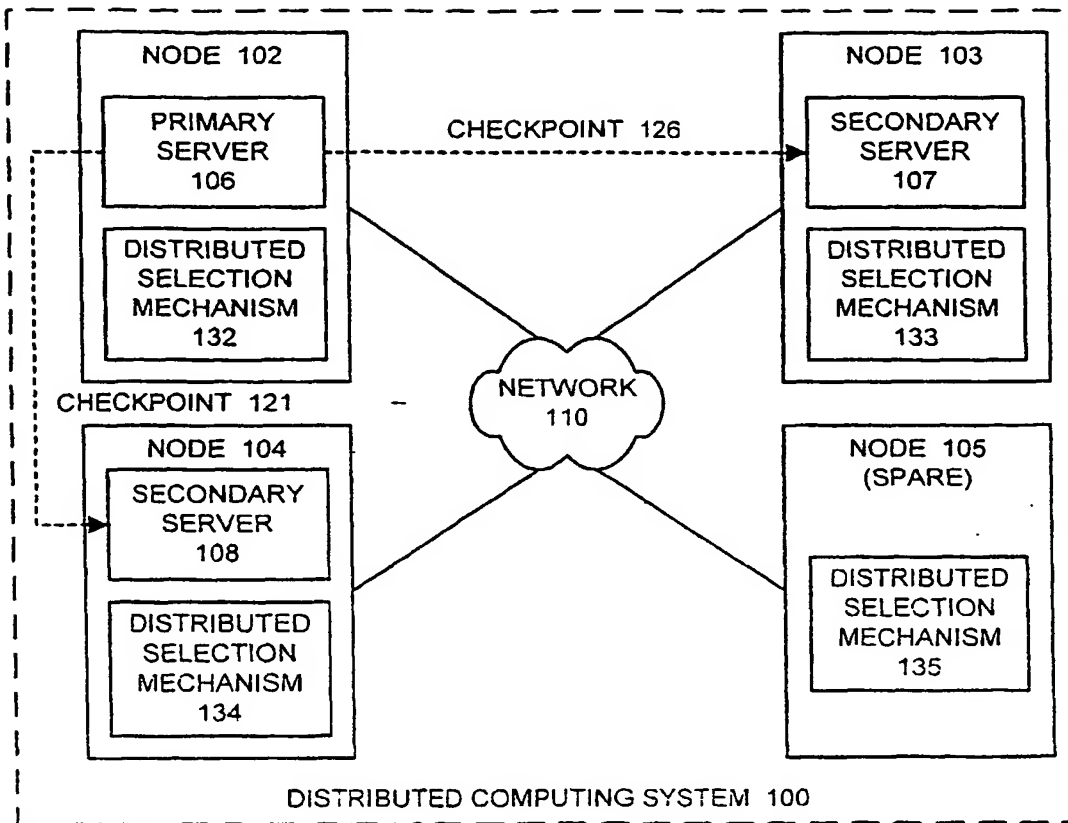


FIG. 1




FIG. 2

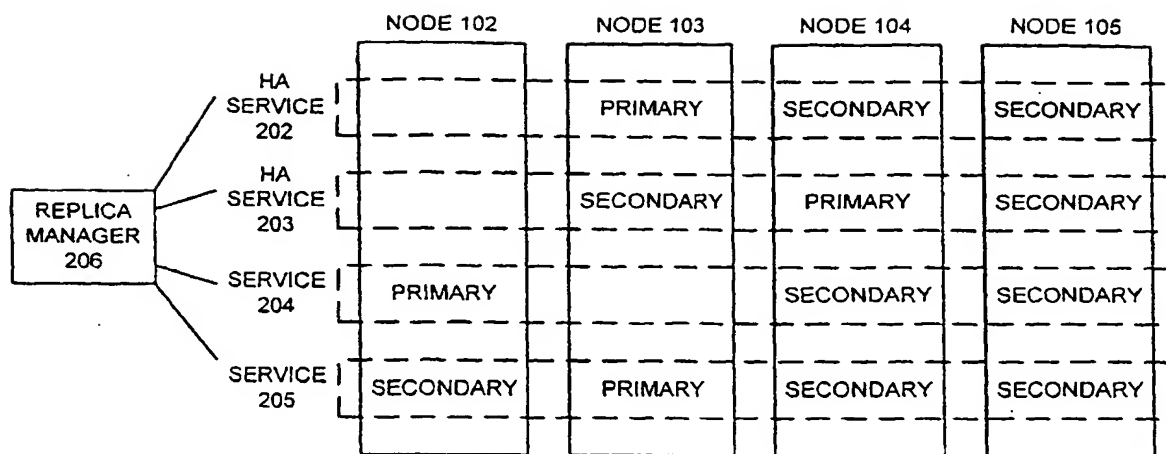


FIG. 3

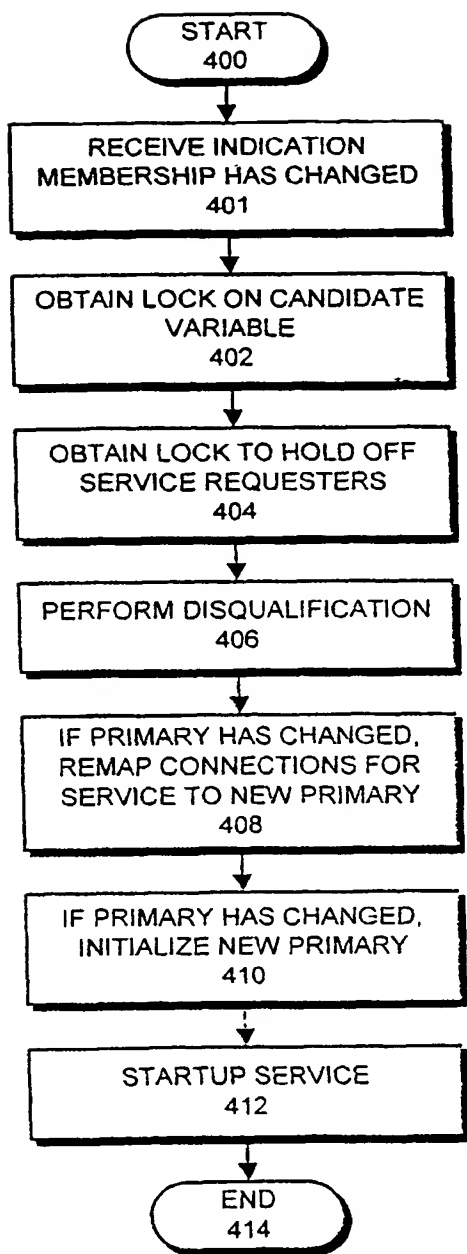


FIG. 4

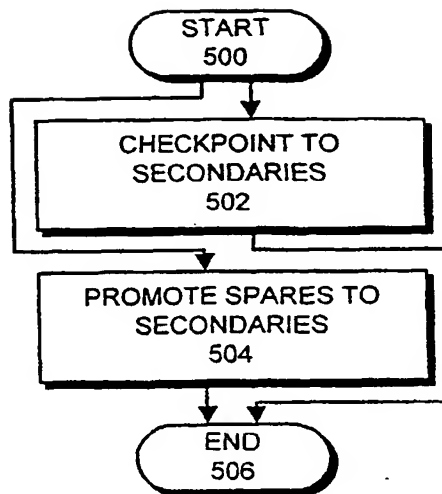


FIG. 5

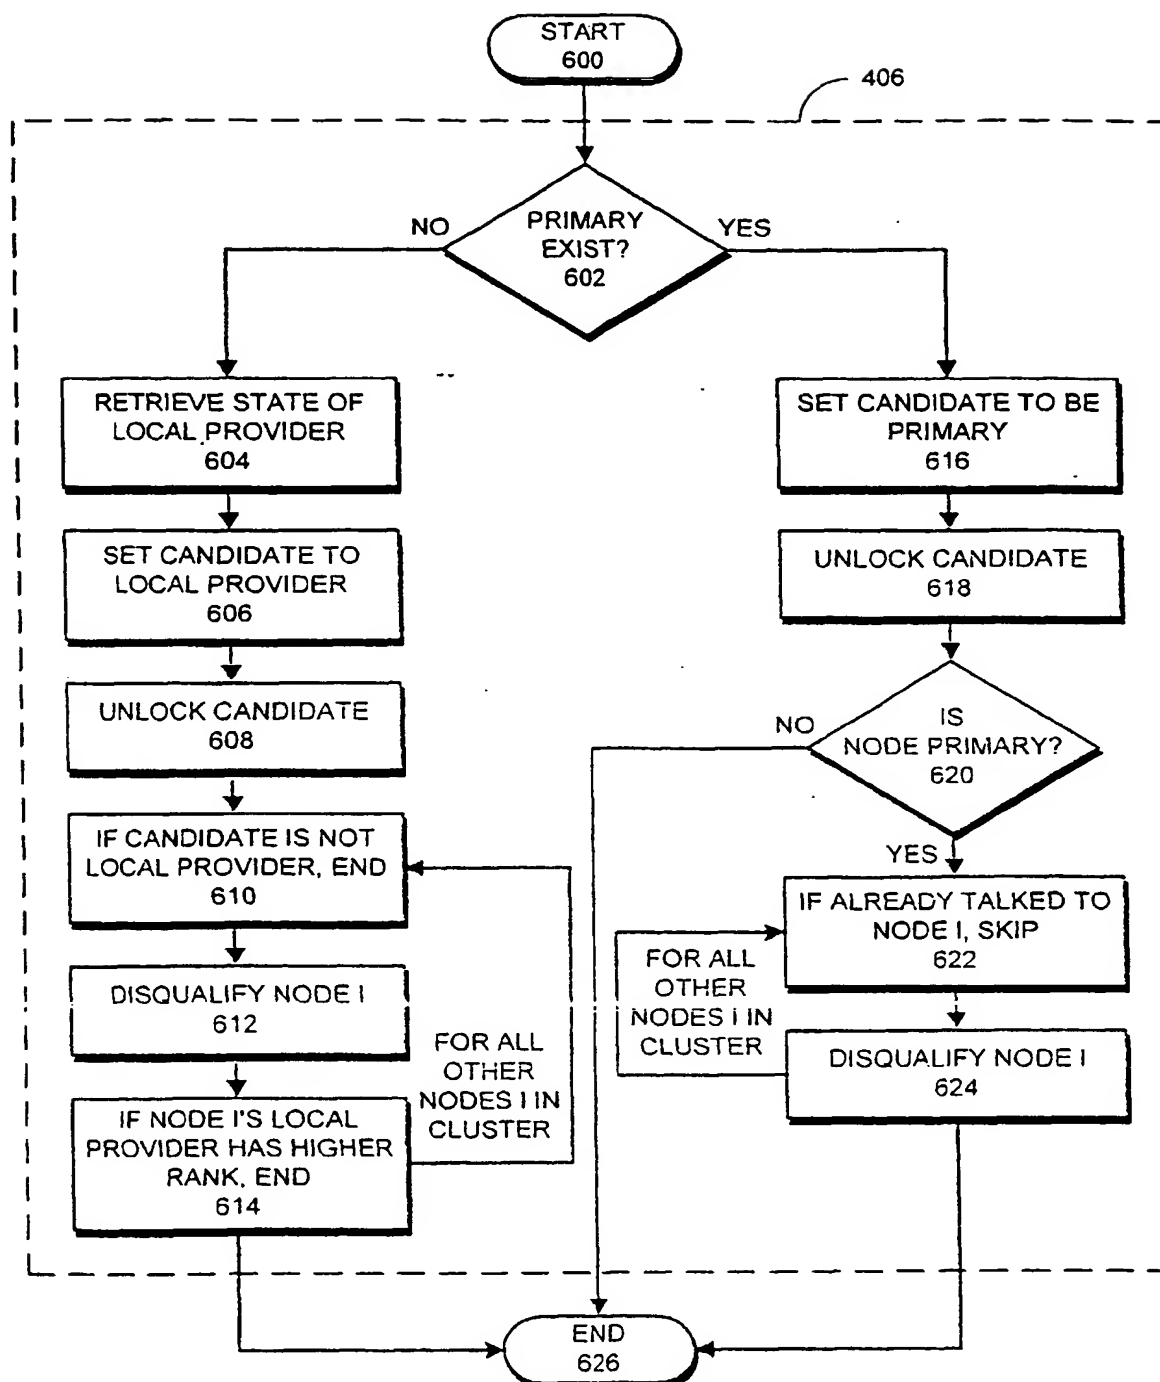


FIG. 6

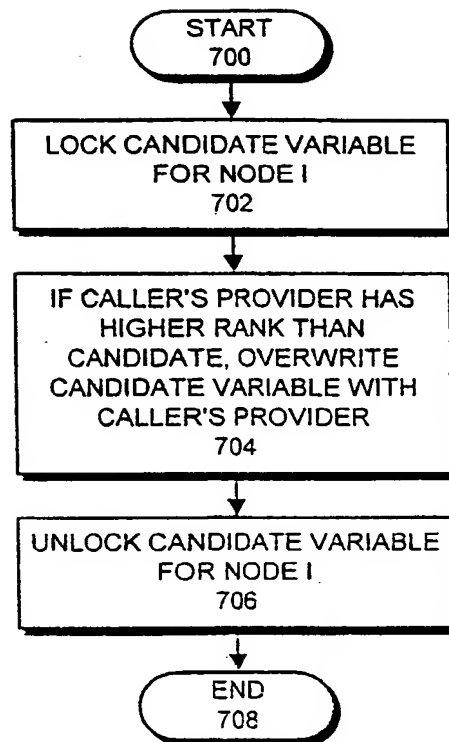



FIG. 7



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) **EP 1 096 751 A3**

(12) **EUROPEAN PATENT APPLICATION**

(88) Date of publication A3:
22.08.2001 Bulletin 2001/34

(51) Int Cl.7: **H04L 29/06**, G06F 9/50,
G06F 11/20

(43) Date of publication A2:
02.05.2001 Bulletin 2001/18

(21) Application number: **00203642.4**

(22) Date of filing: **18.10.2000**

(84) Designated Contracting States:
**AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE**
Designated Extension States:
AL LT LV MK RO SI

(72) Inventors:
• Iterum, Skef F.
San Francisco, CA 94107 (US)
• Murphy, Declan J.
San Francisco, CA 94131 (US)

(30) Priority: 21.10.1999 US 160992 P
15.09.2000 US 662553

(74) Representative: **Hanna, Peter William Derek**
Peter Hanna Associates
11 Mespil Road
Dublin 4 (IE)

(71) Applicant: **SUN MICROSYSTEMS, INC.**
Palo Alto, California 94303 (US)

(54) **Method and apparatus for reaching agreement between nodes in a distributed system**

(57) One embodiment of the present invention provides a system for selecting a node to host a primary server for a service from a plurality of nodes in a distributed computing system. The system operates by receiving an indication that a state of the distributed computing system has changed. In response to this indication, the system determines if there is already a node hosting the primary server for the service. If not, the system selects a node to host the primary server using the assumption that a given node from the plurality of nodes in the dis-

tributed computing system hosts the primary server. The system then communicates rank information between the given node and other nodes in the distributed computing system, wherein each node in the distributed computing system has a unique rank with respect to the other nodes in the distributed computing system. The system next compares the rank of the given node with the rank of the other nodes in the distributed computing system. If one of the other nodes has a higher rank than the given node, the system disqualifies the given node from hosting the primary server.

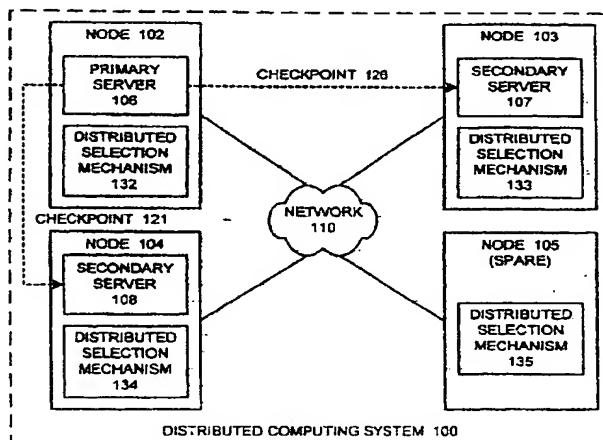


FIG. 1



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 00 20 3642

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	CRISTIAN F ET AL: "FAULT-TOLERANCE IN AIR TRAFFIC CONTROL SYSTEMS" ACM TRANSACTIONS ON COMPUTER SYSTEMS,US,ASSOCIATION FOR COMPUTING MACHINERY. NEW YORK, vol. 14, no. 3, 1 August 1996 (1996-08-01), pages 265-286, XP000638951 ISSN: 0734-2071	1-3, 13-18, 28-32	H04L29/06 G06F9/50 G06F11/20
A	* page 279, line 19 - page 283, last line *	4-12, 19-27	
A	EP 0 750 256 A (DATA GENERAL CORP) 27 December 1996 (1996-12-27)	1,15	
X	* page 7, line 40 - page 11, line 11 *	28	
A	US 5 793 968 A (FARRELL DAVID R ET AL) 11 August 1998 (1998-08-11) * the whole document *	1-32	TECHNICAL FIELDS SEARCHED (Int.Cl.7)
A	CHANG R N ET AL: "A SERVICE ACQUISITION MECHANISM FOR SERVER-BASED HETEROGENEOUS DISTRIBUTED SYSTEMS" IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS,US,IEEE INC, NEW YORK, vol. 5, no. 2, 1 February 1994 (1994-02-01), pages 154-169, XP000440134 ISSN: 1045-9219 * page 158, left-hand column, line 21 - right-hand column, line 19; figure 3 *	1-32	G06F
A	US 5 526 492 A (ISHIDA ATSUNORI) 11 June 1996 (1996-06-11) * column 4, line 39 - column 5, line 17 *	1,15	

	-/--		
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 29 June 2001	Examiner Michel, T
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons</p> <p>& : member of the same patent family, corresponding document</p>			

EPO FORM 1503 03 92 (P4C01)



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 00 20 3642

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
A	WO 97 49039 A (BELL COMMUNICATIONS RES) 24 December 1997 (1997-12-24) * claims 7-10 * * page 11, line 1 * -----	1,7,15, 22,28	
			TECHNICAL FIELDS SEARCHED (Int.Cl.7)
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 29 June 2001	Examiner Michel, T
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			

EPO FORM 1503 03/02 (P04C01)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 00 20 3642

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

29-06-2001

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0750256 A	27-12-1996	US 5666486 A	09-09-1997
		AU 713372 B	02-12-1999
		AU 5601396 A	09-01-1997
		CA 2179473 A	24-12-1996
		JP 9171502 A	30-06-1997
US 5793968 A	11-08-1998	US 5526358 A	11-06-1996
		US 5778185 A	07-07-1998
		AU 3944995 A	22-03-1996
		BR 9508731 A	15-12-1998
		CA 2197324 A	07-03-1996
		CN 1159858 A	17-09-1997
		EP 0776502 A	04-06-1997
		JP 9511115 T	04-11-1997
		WO 9607257 A	07-03-1996
		US 5612957 A	18-03-1997
		US 5699351 A	16-12-1997
US 5526492 A	11-06-1996	JP 4271454 A	28-09-1992
WO 9749039 A	24-12-1997	US 6014686 A	11-01-2000

EPO FORM P0489

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82